


MPAS-Land Ice Model User's Guide

Version: 3.0

Climate, Ocean, Sea-Ice Modeling Team

Los Alamos National Laboratory

September 18, 2013



Foreword

The Model for Prediction Across Scales-Land Ice (MPAS-Land Ice) is an unstructured-mesh land ice model (ice sheets or glaciers) capable of using enhanced horizontal resolution in selected regions of the land ice domain. This allows researchers to perform high-resolution regional simulations at a lower computational cost, while providing realistic ice flow from the low-resolution regions. Model domains may be spherical or on Cartesian domains. MPAS-Land-Ice is being designed for large-scale, hi-resolution simulations of ice sheet dynamics, using a combination of Finite Difference, Finite Volume, and Finite Element Methods on variable resolution meshes. MPAS-Land-Ice will initially be released with support for standard test cases used in verifying model performance. Eventually, support for standard ice sheet configurations will also be included (e.g., stand-alone Greenland and Antarctica simulations).

Prototype dynamical cores for MPAS-Land-Ice have shown good agreement with manufactured solutions and standard test cases, and have been used in land ice evolution experiments aimed at informing the IPCC AR5 on the potential for future sea-level rise from ice sheets (e.g., Ice2Sea international assessment project (Shannon et al., 2013; Edwards et al., 2013)). Two dynamical cores are currently under development for implementation within MPAS-Land-Ice. These include a 1st-order accurate approximation to the momentum balance equations, a prototype of which has been described by Perego et al. (2012), and a "full" Stokes momentum balance, described in Leng et al. (2012).

MPAS-Land Ice is one component within the MPAS framework of climate models that is developed in cooperation between Los Alamos National Laboratory (LANL) and the National Center for Atmospheric Research (NCAR). Functionality that is required by all cores, such as i/o, time management, block decomposition, etc, is developed collaboratively, and this code is shared across cores within the same repository. Each core then solves its own differential equations and physical parameterizations within this framework. This user's guide reflects the spirit of this collaborative process, where Part I, "The MPAS Framework", applies to all cores, and the remaining parts apply to MPAS-Land Ice.

Here we would normally describe the new features of this version. For the initial release, we will simply review the major features of the basic MPAS-Land Ice model. We employ a finite-volume discretization of the ice continuity equation using a C-grid staggering in the horizontal. The vertical coordinate is sigma. The time-stepping method is Forward Euler (explicit). Ice advection is performed by first-order upwinding. No tracer advection is available at present. In the initial release velocity can only be solved using the Shallow Ice Approximation.

A history of past releases of the Land Ice core within the MPAS version numbering scheme is as follows:

version	date	description
2.0.0	November 15, 2013	Initial public release of Land Ice core (SIA velocity solver only)
3.0.0	November 18, 2013	Fix bug in SIA slope calculation. Introduction of run-time I/O streams.

Information about MPAS-Land Ice, including the most recent code, user's guide, and test cases,

may be found at <http://mpas-dev.github.com>. This user's guide refers to version 3.0.

Contributors to this guide:

Matt Hoffman, Stephen Price

Additional contributors to MPAS Framework sections:

Michael Duda, Douglas Jacobsen

Funding for the development of MPAS-Land Ice was provided by the United States Department of Energy, Office of Science.

Contents

1	MPAS-Land Ice Quick Start Guide	8
I	The MPAS Framework	9
2	Building MPAS	10
2.1	Prequisites	10
2.2	Compiling I/O Libraries	10
2.2.1	netCDF	10
2.2.2	parallel-netCDF	11
2.2.3	PIO	11
2.3	Compiling MPAS	11
2.4	Cleaning	13
2.5	Graph partitioning with METIS	13
3	Grid Description	15
4	Configuring Model Input and Output	19
4.1	XML stream configuration files	19
4.2	Optional stream attributes	21
4.3	Stream definition examples	22
4.3.1	Example: a single-precision output stream with one month of data per file	22
4.3.2	Example: appending records to existing output files	23
4.3.3	Example: referencing filename intervals to a time other than the start time	24
5	Visualization	26
5.1	ParaView	26
II	MPAS-Land Ice	29
6	Governing Equations	30
6.1	Momentum Balance	30
6.2	Time Integration	31
6.3	Advection	31

7	Model Configuration	32
7.1	Dimensions	32
7.2	Namelist options	32
7.2.1	velocity_solver	32
7.2.2	advection	33
7.2.3	physical_parameters	33
7.2.4	time_integration	33
7.2.5	time_management	34
7.2.6	io	34
7.2.7	decomposition	35
7.2.8	debug	35
7.3	Variable definitions	36
7.3.1	state	36
7.3.2	tend	36
7.3.3	mesh	37
7.4	Run-time input/output streams	38
7.4.1	input	38
7.4.2	output	38
7.4.3	restart	39
7.4.4	basicmesh	40
7.4.5	Other streams	40
8	Land Ice Visualization	41
8.1	Python	41
9	Test Cases	42
9.1	Halfar Dome	42
9.1.1	Provided Files	42
9.1.2	Results	43
9.2	EISMINT-1 Test Cases	45
9.2.1	Provided Files	45
9.2.2	Results	45
9.3	Real World Test Cases	46
10	Global Statistics	47
11	Running MPAS-Land Ice within a coupled climate model	48
12	Troubleshooting	49
12.1	Choice of time step	49
13	Known Issues	50

III Bibliography 51

IV Appendices 53

A Namelist options 54

A.1	config_block_decomp_file_prefix	54
A.2	config_calendar_type	54
A.3	config_default_flowParamA	54
A.4	config_do_restart	55
A.5	config_dt	55
A.6	config_dynamic_thickness	55
A.7	config_explicit_proc_decomp	56
A.8	config_flowLawExponent	56
A.9	config_ice_density	56
A.10	config_num_halos	56
A.11	config_number_of_blocks	57
A.12	config_ocean_density	57
A.13	config_pio_num_iotasks	57
A.14	config_pio_stride	58
A.15	config_print_thickness_advection_info	58
A.16	config_proc_decomp_file_prefix	58
A.17	config_restart_timestamp_name	58
A.18	config_run_duration	59
A.19	config_sea_level	59
A.20	config_start_time	59
A.21	config_stop_time	60
A.22	config_thickness_advection	60
A.23	config_time_integration	60
A.24	config_tracer_advection	61
A.25	config_velocity_solver	61
A.26	config_write_output_on_startup	61
A.27	config_year_digits	61

B Variable definitions 63

B.1	mesh	63
B.1.1	angleEdge	63
B.1.2	areaCell	63
B.1.3	areaTriangle	63
B.1.4	bedTopography	64
B.1.5	cellTangentPlane	64
B.1.6	cellsOnCell	64
B.1.7	cellsOnEdge	65
B.1.8	cellsOnVertex	65
B.1.9	coeffs_reconstruct	65
B.1.10	dcEdge	66
B.1.11	dvEdge	66
B.1.12	edgeNormalVectors	66

B.1.13	edgeSignOnCell	66
B.1.14	edgeSignOnVertex	67
B.1.15	edgesOnCell	67
B.1.16	edgesOnEdge	67
B.1.17	edgesOnVertex	68
B.1.18	indexToCellID	68
B.1.19	indexToEdgeID	68
B.1.20	indexToVertexID	69
B.1.21	kiteAreasOnVertex	69
B.1.22	latCell	69
B.1.23	latEdge	70
B.1.24	latVertex	70
B.1.25	layerCenterSigma	70
B.1.26	layerInterfaceSigma	70
B.1.27	layerThicknessFractions	71
B.1.28	localVerticalUnitVectors	71
B.1.29	lonCell	71
B.1.30	lonEdge	72
B.1.31	lonVertex	72
B.1.32	nEdgesOnCell	72
B.1.33	nEdgesOnEdge	73
B.1.34	sfcMassBal	73
B.1.35	verticesOnCell	73
B.1.36	verticesOnEdge	73
B.1.37	weightsOnEdge	74
B.1.38	xCell	74
B.1.39	xEdge	74
B.1.40	xVertex	75
B.1.41	yCell	75
B.1.42	yEdge	75
B.1.43	yVertex	76
B.1.44	zCell	76
B.1.45	zEdge	76
B.1.46	zVertex	77
B.2	state	77
B.2.1	cellMask	77
B.2.2	edgeMask	77
B.2.3	layerThickness	78
B.2.4	layerThicknessEdge	78
B.2.5	lowerSurface	78
B.2.6	normalVelocity	78
B.2.7	temperature	79
B.2.8	thickness	79
B.2.9	uReconstructMeridional	79
B.2.10	uReconstructX	80
B.2.11	uReconstructY	80
B.2.12	uReconstructZ	80
B.2.13	uReconstructZonal	81

B.2.14	upperSurface	81
B.2.15	upperSurfaceVertex	81
B.2.16	vertexMask	81
B.2.17	xtime	82
B.3	tend	82
B.3.1	tend_layerThickness	82
B.3.2	tend_temperature	82

Chapter 1

MPAS-Land Ice Quick Start Guide

This chapter provides MPAS-Land Ice users with a quick start description of how to build and run the model. It is meant merely as a brief overview of the process, while the more detailed descriptions of each step are provided in later sections.

In general, the build process follows the following steps.

1. Build MPI Layer (OpenMPI, MVAPICH2, etc.)
2. Build serial NetCDF library (v3.6.3, v4.1.3, etc.)
3. Build Parallel-NetCDF library (v1.2.1, v1.3.0, etc.)
4. Build Parallel I/O library (v1.4.1, v1.6.1, etc.)
5. (Optional) Build METIS library and executables (v4.0, v5.0.2, etc.)
6. Checkout MPAS-Land Ice from repository
7. Build Land Ice core

After step 7, an executable should be created called `landice_model.exe`. Once the executable is built, one can begin the run process as follows:

1. Create run directory.
2. Copy executable to run directory.
3. Copy `namelist.input` into run directory.
4. (Optional) Copy input and graph files into run directory.
5. Edit `namelist.input` to have the proper parameters.
If step 4 was skipped, ensure paths to input and graph files are appropriately set.
6. (Optional) Create graph files, using METIS executable (`pmetis` or `gpmets` depending on version).
A graph file is required for each processor count you want to use.
7. Run MPAS-Land Ice.
8. Visualize output file, and perform analysis.

Part I

The MPAS Framework

Chapter 2

Building MPAS

2.1 Prerequisites

To build MPAS, compatible C and Fortran compilers are required. Additionally, the MPAS software relies on the PIO parallel I/O library to read and write model fields, and the PIO library requires the standard netCDF library as well as the parallel-netCDF library from Argonne National Labs. All libraries must be compiled with the same compilers that will be used to build MPAS. Section 2.2 summarizes the basic procedure of installing the required I/O libraries for MPAS.

In order for the MPAS makefiles to find the PIO, parallel-netCDF, and netCDF include files and libraries, the environment variables `PIO`, `PNETCDF`, and `NETCDF` should be set to the root installation directories of the PIO, parallel-netCDF, and netCDF installations, respectively. Newer versions of the netCDF library use a separate Fortran interface library; the top-level MPAS Makefile attempts to add `-lnetcdf` to the linker flags, but some linkers require that `-lnetcdf` appear before `-lnetcdf`, in which case `-lnetcdf` will need to be manually added just before `-lnetcdf` in the specification of `LIBS` in the top-level Makefile.

An MPI installation such as MPICH or OpenMPI is also required, and there is no option to build a serial version of the MPAS executables. There is currently no support for shared-memory parallelism with OpenMP within the MPAS framework.

2.2 Compiling I/O Libraries

NOTE: It's important to note the MPAS Developers are not responsible for any of the libraries that are used within MPAS. Support for specific libraries should be taken up with the respective developer groups.

Although most recent versions of the I/O libraries should work, the most tested versions of these libraries are: netCDF 4.1.3, parallel-netCDF 1.3.1, and PIO 1.4.1. The netCDF and parallel-netCDF libraries must be installed before building PIO library.

All commands are presented for `csh`, and will not work if pasted into another shell. Please translate them to the appropriate commands in your shell.

2.2.1 netCDF

Version 4.1.3 of the netCDF library may be downloaded from http://www.unidata.ucar.edu/downloads/netcdf/netcdf-4_1_3/index.jsp. Assuming the `gfortran` and `gcc` compilers will be used, the following shell commands are generally sufficient to install netCDF.

```

> setenv FC gfortran
> setenv F77 gfortran
> setenv F90 gfortran
> setenv CC gcc
> ./configure --prefix=XXXXX --disable-dap --disable-netcdf-4 --disable-cxx
--disable-shared --enable-fortran
> make all check
> make install

```

Here, XXXXX should be replaced with the directory that will serve as the root installation directory for netCDF. *Before proceeding to compile PIO the NETCDF_PATH environment variable should be set to the netCDF root installation directory.*

Certain compilers require addition flags in the CPPFLAGS environment variable. Please refer to the netCDF installation instructions for these flags.

2.2.2 parallel-netCDF

Version 1.3.1 of the parallel-netCDF library may be downloaded from <https://trac.mcs.anl.gov/projects/parallel-netcdf/wiki/Download>. Assuming the gfortran and gcc compilers will be used, the following shell commands are generally sufficient to install parallel-netCDF.

```

> setenv MPIF90 mpif90
> setenv MPIF77 mpif90
> setenv MPICC mpicc
> ./configure --prefix=XXXXX
> make
> make install

```

Here, XXXXX should be replaced with the directory that will serve as the root installation directory for parallel-netCDF. *Before proceeding to compile PIO the PNETCDF_PATH environment variable should be set to the parallel-netCDF root installation directory.*

2.2.3 PIO

Instructions for building PIO can be found at <http://www.cesm.ucar.edu/models/pio/>. Please refer to these instructions for building PIO.

After PIO is built, and installed the PIO environment variable needs to be defined to point at the directory PIO is installed into. Older versions of PIO cannot be installed, and the PIO environment variable needs to be set to the directory where PIO was built instead.

2.3 Compiling MPAS

Before compiling MPAS, the NETCDF, PNETCDF, and PIO environment variables must be set to the library installation directories as described in the previous section. A CORE variable also needs to either be defined or passed in during the make process. If CORE is not specified, the build process will fail.

The MPAS code uses only the ‘make’ utility for compilation. Rather than employing a separate configuration step before building the code, all information about compilers, compiler flags, etc.,

is contained in the top-level `Makefile`; each supported combination of compilers (i.e., a configuration) is included in the `Makefile` as a separate make target, and the user selects among these configurations by running `make` with the name of a build target specified on the command-line, e.g.,

```
> make gfortran
```

to build the code using the GNU Fortran and C compilers. Some of the available targets are listed in the table below, and additional targets can be added by simply editing the `Makefile` in the top-level directory.

Target	Fortran compiler	C compiler	MPI wrappers
<code>xlf</code>	<code>xlf90</code>	<code>xlc</code>	<code>mpxlf90 / mpcc</code>
<code>pgi</code>	<code>pgf90</code>	<code>pgcc</code>	<code>mpif90 / mpicc</code>
<code>ifort</code>	<code>ifort</code>	<code>gcc</code>	<code>mpif90 / mpicc</code>
<code>gfortran</code>	<code>gfortran</code>	<code>gcc</code>	<code>mpif90 / mpicc</code>
<code>g95</code>	<code>g95</code>	<code>gcc</code>	<code>mpif90 / mpicc</code>

In order to get a more complete and up-to-date list of available targets, one can use the following command within the top-level of MPAS. **NOTE:** This command is known to not work with Mac OSX.

```
> make -rpn | sed -n -e '/^$/ { n ; /^[^]*:/p }' | sed "s/: *.*$/g"
```

The MPAS framework supports multiple *cores* — currently a shallow water model, an ocean model, a non-hydrostatic atmosphere model, a non-hydrostatic atmosphere initialization core, and a land ice core — so the build process must be told which core to build. This is done by either setting the environment variable `CORE` to the name of the model core to build, or by specifying the core to be built explicitly on the command-line when running `make`. For the shallow water core, for example, one may run either

```
> setenv CORE sw
> make gfortran
```

or

```
> make gfortran CORE=sw
```

If the `CORE` environment variable is set and a core is specified on the command-line, the command-line value takes precedence; if no core is specified, either on the command line or via the `CORE` environment variable, the build process will stop with an error message stating such. Assuming compilation is successful, the model executable, named `_${CORE}_model` (e.g., `sw_model`), should be created in the top-level MPAS directory.

In order to get a list of available cores, one can simply run the top-level `Makefile` without setting the `CORE` environment variable, or passing the core via the command-line. An example of the output from this can be seen below.

```
> make
```

```
( make error )
make[1]: Entering directory 'mpas'

Usage: make target CORE=[core] [options]

Example targets:
ifort
gfortran
xlf
pgi

Available Cores:
atmosphere
init_atmosphere
landice
ocean
sw

Available Options:
DEBUG=true      - builds debug version. Default is optimized version.
USE_PAPI=true   - builds version using PAPI for timers. Default is off.
TAU=true        - builds version using TAU hooks for profiling. Default is off.

Ensure that NETCDF, PNETCDF, PIO, and PAPI (if USE_PAPI=true) are environment variables
that point to the absolute paths for the libraries.

***** ERROR *****
No CORE specified. Quitting.
***** ERROR *****

make[1]: Leaving directory 'mpas'
```

2.4 Cleaning

To remove all files that were created when the model was built, including the model executable itself, `make` may be run for the 'clean' target:

```
> make clean
```

As with compiling, the core to be cleaned is specified by the `CORE` environment variable, or by specifying a core explicitly on the command-line with `CORE=`.

2.5 Graph partitioning with METIS

Before MPAS can be run in parallel, a mesh decomposition file with an appropriate number of partitions (equal to the number of MPI tasks that will be used) is required in the run directory. A limited number of mesh decomposition files (`graph.info.part.*`) are provided with each test case. In order to create new mesh decomposition files for your desired number of partitions, begin with the provided `graph.info` file and partition with METIS software (<http://glaros.dtc.umn.edu/gkhome/views/metis>). The serial graph partitioning program, METIS (rather than ParMETIS or

hMETIS) should be sufficient for quickly partitioning any SCVT produced by the `grid_gen` mesh generator.

After installing METIS, a `graph.info` file may be partitioned into N partitions by running

```
> gpmetis graph.info N
```

The resulting file, `graph.info.part.N`, can then be copied into the MPAS run directory before running the model with N MPI tasks.

Chapter 3

Grid Description

This chapter provides a brief introduction to the common types of grids used in the MPAS framework.

The MPAS grid system requires the definition of seven elements. These seven elements are composed of two types of *cells*, two types of *lines*, and three types of *points*. These elements are depicted in Figure 3.1 and defined in Table 3.1. These elements can be defined on either the plane or the surface of the sphere. The two types of cells form two meshes, a primal mesh composed of Voronoi regions and a dual mesh composed of Delaunay triangles. Each corner of a primal mesh cell is uniquely associated with the “center” of a dual mesh cell and vice versa. So we define the two mesh as either a primal mesh (composed of cells P_i) or a dual mesh (composed of cells D_v). The center of any primal mesh cell, P_i , is denoted by \mathbf{x}_i and the center of any the dual mesh cell, D_v , is denoted by \mathbf{x}_v . The boundary of a given primal mesh cell P_i is composed of the set of lines that connect the \mathbf{x}_v locations of associated dual mesh cells D_v . Similarly, the boundary of a given dual mesh cell D_v is composed of the set of lines that connect the \mathbf{x}_i locations of the associated primal mesh cells P_i .

As shown in Figure 3.1, a line segment that connects two primal mesh cell centers is uniquely associated with a line segment that connects two dual mesh cell centers. We assume that these two line segments cross and the point of intersection is labeled as \mathbf{x}_e . In addition, we assume that these two line segments are orthogonal as indicated in Figure 3.1. Each \mathbf{x}_e is associated with two distances: d_e measures the distance between the primal mesh cells sharing \mathbf{x}_e and l_e measures the distance between the dual mesh cells sharing \mathbf{x}_e .

Since the two line segments crossing at \mathbf{x}_e are orthogonal, these line segments form a convenient local coordinate system for each edge. At each \mathbf{x}_e location a unit vector \mathbf{n}_e is defined to be parallel to the line connecting primal mesh cells. A second unit vector \mathbf{t}_e is defined such that $\mathbf{t}_e = \mathbf{k} \times \mathbf{n}_e$.

In addition to these seven element types, we require the definition of *sets of elements*. In all, eight different types of sets are required and these are defined and explained in Table 3.2 and Figure 3.2. The notation is always of the form of, for example, $i \in CE(e)$, where the LHS indicates the type of element to be gathered (cells) based on the RHS relation to another type of element (edges).

Table 3.3 provides the names of all *elements* and all *sets of elements* as used in the MPAS framework. Elements appear twice in the table when described in the grid file in more than one way, e.g. points are described with both cartesian and latitude/longitude coordinates. An “ncdump -h” of any MPAS grid, output or restart file will contain all variable names shown in second column of Table 3.3.

Table 3.1: Definition of elements used to build the MPAS grid.

<i>Element</i>	<i>Type</i>	<i>Definition</i>
\mathbf{x}_i	point	location of center of primal-mesh cells
\mathbf{x}_v	point	location of center of dual-mesh cells
\mathbf{x}_e	point	location of edge points where velocity is defined
d_e	line segment	distance between neighboring \mathbf{x}_i locations
l_e	line segment	distance between neighboring \mathbf{x}_v locations
P_i	cell	a cell on the primal-mesh
D_v	cell	a cell on the dual-mesh

Table 3.2: Definition of element groups used to reference connections in the MPAS grid.
Examples are provided in Figure 3.2.

<i>Syntax</i>	<i>output</i>
$e \in EC(i)$	set of edges that define the boundary of P_i .
$e \in EV(v)$	set of edges that define the boundary of D_v .
$i \in CE(e)$	two primal-mesh cells that share edge e .
$i \in CV(v)$	set of primal-mesh cells that form the vertices of dual mesh cell D_v .
$v \in VE(e)$	the two dual-mesh cells that share edge e .
$v \in VI(i)$	the set of dual-mesh cells that form the vertices of primal-mesh cell P_i .
$e \in ECP(e)$	edges of cell pair meeting at edge e .
$e \in EVC(v, i)$	edge pair associated with vertex v and mesh cell i .

Table 3.3: Variable names used to describe a MPAS grid.

<i>Element</i>	<i>Name</i>	<i>Size</i>	<i>Comment</i>
\mathbf{x}_i	{x,y,z}Cell	nCells	cartesian location of \mathbf{x}_i
\mathbf{x}_i	{lon,lat}Cell	nCells	longitude and latitude of \mathbf{x}_i
\mathbf{x}_v	{x,y,z}Vertex	nVertices	cartesian location of \mathbf{x}_v
\mathbf{x}_v	{lon,lat}Vertex	nVertices	longitude and latitude of \mathbf{x}_v
\mathbf{x}_e	{x,y,z}Edge	nEdges	cartesian location of \mathbf{x}_e
\mathbf{x}_e	{lon,lat}Edge	nEdges	longitude and latitude of \mathbf{x}_e
d_e	dcEdge	nEdges	distance between \mathbf{x}_i locations
l_e	dvEdge	nEdges	distance between \mathbf{x}_v locations
$e \in EC(i)$	edgesOnCell	(nEdgesMax,nCells)	edges that define P_i .
$e \in EV(v)$	edgesOnVertex	(3,nCells)	edges that define D_v .
$i \in CE(e)$	cellsOnEdge	(2,nEdges)	primal-mesh cells that share edge e .
$i \in CV(v)$	cellsOnVertex	(3,nVertices)	primal-mesh cells that define D_v .
$v \in VE(e)$	verticesOnEdge	(2,nEdges)	dual-mesh cells that share edge e .
$v \in VI(i)$	verticesOnCell	(nEdgesMax,nCells)	vertices that define P_i .




Figure 3.1: Definition of elements used to build the MPAS grid. Also see Table 3.1.




Figure 3.2: Definition of element groups used to reference connections in the MPAS grid.
Also see Table 3.2.

Chapter 4

Configuring Model Input and Output

The reading and writing of model fields in MPAS is handled by user-configurable *streams*. A stream represents a fixed set of model fields, together with dimensions and attributes, that are all written or read together to or from the same file or set of files. Each MPAS model core may define its own set of default streams that it typically uses for reading initial conditions, for writing and reading restart fields, and for writing additional model history fields. Besides these default streams, users may define new streams to, e.g., write certain diagnostic fields at a higher temporal frequency than the usual model history fields.

Streams are defined in XML configuration files that are created at build time for each model core. The name of this XML file is simply ‘streams.’ suffixed with the name of the core. For example, the streams for the *sw* (shallow-water) core are defined in a file named ‘streams.sw’. An XML stream file may further reference other text files that contain lists of the model fields that are read or written in each of the streams defined in the XML stream file.

Changes to the XML stream configuration file will take effect the next time an MPAS core is run; there is no need to re-compile after making modifications to the XML files. As described in the next section, it is therefore possible, e.g., to change the interval at which a stream is written, the template for the filenames associated with a stream, or the set of fields that are written to a stream, without the need to re-compile any code.

Two classes of streams exist in MPAS: *immutable* streams and *mutable* streams. Immutable streams are those for which the set of fields that belong to the stream may not be modified at model run-time; however, it is possible to modify the interval at which the stream is read or written, the filename template describing the files containing the stream on disk, and several other parameters of the stream. In contrast, all aspects of mutable streams, including the set of fields that belong to the stream, may be modified at run-time. The motivation for the creation of two stream classes is the idea that an MPAS core may not function correctly if certain fields are not read in upon model start-up or written to restart files, and it is therefore not reasonable for users to modify this set of required fields at run-time. An MPAS core developer may choose to implement such streams as immutable streams. Since fields may not be added to an immutable stream at run-time, new immutable streams may not be defined at run-time, and the only type of new stream that may be defined at run-time is the mutable stream type.

4.1 XML stream configuration files

The XML stream configuration file for an MPAS core always has a parent XML *element* named **streams**, within which individual streams are defined:

<streams>

... one or more stream definitions ...

</streams>

Immutable streams are defined with the `immutable_stream` element, and mutable streams are defined with the `stream` element:

```
<immutable_stream name="initial_conditions"
                  type="input"
                  filename_template="init.nc"
                  input_interval="initial_only"
                  />
```

```
<stream name="history"
         type="output"
         filename_template="output.$Y-$M-$D.$h.$m.$s.nc"
         output_interval="6:00:00" />
```

... model fields belonging to this stream ...

</stream>

As shown in the example stream definitions, above, both classes of stream have the following required attributes:

- **name** — A unique name used to refer to the stream.
- **type** — The type of stream, either "input", "output", "input;output", or "none". A stream may be both an input and an output stream (i.e., "input;output") if, for example, it is read once at model start-up to provide initial conditions and thereafter written periodically to provide model checkpoints. A stream may be defined as neither input nor output (i.e., "none") for the purposes of defining a set of fields for inclusion other streams. Note that, for immutable streams, the type attribute may not be changed at run-time.
- **filename_template** — The template for files that exist or will be created by the stream. The filename template may include any of the following variables, which are expanded based on the simulated time at which files are first created.
 - **\$Y** — Year
 - **\$M** — Month
 - **\$D** — Day of the month
 - **\$d** — Day of the year
 - **\$h** — Hour

- **\$m** — Minute
- **\$s** — Second

A filename template may include either a relative or an absolute path, in which case MPAS will attempt to create any directories in the path that do not exist, subject to filesystem permissions.

- **input_interval** — For streams that have type "input" or "input;output", the interval, beginning at the model initial time, at which the stream will be read. Possible values include a time interval specification in the format "YYYY-MM-DD_hh:mm:ss"; the value "initial_only", which specifies that the stream is read only once at the model initial time; or the value "none", which specifies that the stream is not read during a model run.
- **output_interval** — For streams that have type "output" or "input;output", the interval, beginning at the model initial time, at which the stream will be written. Possible values include a time interval specification in the format "YYYY-MM-DD_hh:mm:ss"; the value "initial_only", which specifies that the stream is written only once at the model initial time; or the value "none", which specifies that the stream is not written during a model run.

Finally, the set of fields that belong to a mutable stream may be specified with any combination of the following elements. Note that, for immutable streams, no fields are specified at run-time in the XML configuration file.

- **var** — Associates the specified variable with the stream. The variable may be any of those defined in an MPAS core's Registry.xml file, but may not include individual constituent arrays from a var_array.
- **var_array** — Associates all constituent variables in a var_array, defined in an MPAS core's Registry.xml file, with the stream.
- **var_struct** — Associates all variables in a var_struct, defined in an MPAS core's Registry.xml file, with the stream.
- **stream** — Associates all explicitly associated fields in the specified stream with the stream; streams are not recursively included.
- **file** — Associates all variables listed in the specified text file, with one field per line, with the stream.

4.2 Optional stream attributes

Besides the required attributes described in the preceding section, several additional, optional attributes may be added to the definition of a stream.

- **filename_interval** — The interval between the timestamps used in the construction of the names of files associated with a stream. Possible values include a time interval specification in the format "YYYY-MM-DD_hh:mm:ss"; the value "none", indicating that only one file containing all times is associated with the stream; the value "input_interval" that, for input type streams, indicates that each time to be read from the stream will come from a unique file; or the value "output_interval" that, for output type streams, indicates that each time to

be written to the stream will go to a unique file whose name is based on the timestamp of the data being written. The default value is `"input_interval"` for input type streams and `"output_interval"` for output type streams. For streams of type `"input;output"`, the default filename interval is `"input_interval"` if the input interval is an interval (i.e., not `"initial_only"`), or `"output_interval"` otherwise. Refer to Section 4.3.1 for an example of the use of the `filename_interval` attribute.

- **reference_time** — A time that is an integral number of filename intervals from the timestamp of any file associated with the stream. The default value is the start time of the model simulation. Refer to Section 4.3.3 for an example of the use of the `reference_time` attribute.
- **clobber_mode** — Specifies how a stream should handle attempts to write to a file that already exists. Possible values for the mode include:
 - **"overwrite"** — The stream is allowed to overwrite records in existing files and to append new records to existing files; records not explicitly written to are left untouched.
 - **"truncate"** or **"replace_files"** — The stream is allowed to overwrite existing files, which are first truncated to remove any existing records; this is equivalent to replacing any existing files with newly created files of the same name.
 - **"append"** — The stream is only allowed to append new records to existing files; existing records may not be overwritten.
 - **"never_modify"** — The stream is not allowed to modify existing files in any way.

The default clobber mode for streams is `"never_modify"`. Refer to Section 4.3.2 for an example of the use of the `clobber_mode` attribute.

- **precision** — The precision with which real-valued fields will be written or read in a stream. Possible values include `"single"` for 4-byte real values, `"double"` for 8-byte real values, or `"native"`, which specifies that real-valued fields will be written or read in whatever precision the MPAS core was compiled. The default value is `"native"`. Refer to Section 4.3.1 for an example of the use of the `precision` attribute.
- **packages** — A list of packages attached to the stream. A stream will be active (i.e., read or written) only if at least one of the packages attached to it is active, or if no packages at all are attached. Package names are provided as a semi-colon-separated list. Note that packages may only be defined in an MPAS core's `Registry.xml` file at build time. By default, no packages are attached to a stream.

4.3 Stream definition examples

This section provides several example streams that make use of the optional stream attributes described in Section 4.2. All examples are of output streams, since it is more likely that a user will need to write additional fields than to read additional fields, which a model would need to be aware of; however, the concepts that are illustrated here translate directly to input streams as well.

4.3.1 Example: a single-precision output stream with one month of data per file

In this example, the optional attribute specification `filename_interval="01-00.00:00:00"` is added to force a new output file to be created for the stream every month. Note that the general

format for time interval specifications is YYYY-MM-DD.hh:mm:ss, where any leading terms can be omitted; in this case, the year part of the interval is omitted. To reduce the file size, the specification `precision="single"` is also added to force real-valued fields to be written as 4-byte floating-point values, rather than the default of 8 bytes.

```
<stream name="diagnostics"
      type="output"
      filename_template="diagnostics.$Y-$M.nc"
      filename_interval="01-00_00:00:00"
      precision="single"
      output_interval="6:00:00" />

      <var name="u10"/>
      <var name="v10"/>
      <var name="t2"/>
      <var name="q2"/>

</stream>
```

The only fields that will be written to this stream are the hypothetical 10-m diagnosed wind components, the 2-m temperature, and the 2-m specific humidity variables. Also, note that the filename template only includes the year and month from the model valid time; this can be problematic when the simulation starts in the middle of a month, and a solution for this problem is illustrated in the example of Section [4.3.3](#).

4.3.2 Example: appending records to existing output files

By default, streams will never modify existing files whose filenames match the name of a file that would otherwise be written during the course of a simulation. However, when restarting a simulation that is expected to add more records to existing output files, it can be useful to instruct the MPAS I/O system to append these records, thereby modifying existing files. This may be accomplished with the `clobber_mode` attribute.

```
<stream name="diagnostics"
      type="output"
      filename_template="diagnostics.$Y-$M.nc"
      filename_interval="01-00_00:00:00"
      precision="single"
      clobber_mode="append"
      output_interval="6:00:00" />

      <var name="u10"/>
      <var name="v10"/>
      <var name="t2"/>
      <var name="q2"/>

</stream>
```


In general, if MPAS were to attempt to write a record at a time that already existed in an output file, a clobber_mode of ‘append’ would not permit the write to take place, since this would modify existing data; in ‘append’ mode, only new records may be added. However, due to a peculiarity in the implementation of the ‘append’ clobber mode, it may be possible for an output file to contain duplicate times. This can happen when the first record that is appended to an existing file has a timestamp not matching any in the file, after which, any record that is written — regardless of whether its timestamp matches one already in the file — will be appended to the end of the file. This situation may arise, for example, when restarting a model simulation with a shorter output_interval than was used in the original model simulation with an MPAS core that does not write the first output time for restart runs.

4.3.3 Example: referencing filename intervals to a time other than the start time

The example stream of the previous sections creates a new file each month during the simulation, and the filenames contain only the year and month of the timestamp when the file was created. If a simulation begins at 00 UTC on the first day of a month, then each file in the diagnostic stream will contain only output times that fall within the month in the filename. However, if a simulation were to begin in the middle of a month — for example, the month of June, 2014 — the first diagnostics output file would have a filename of ‘diagnostics.2014-06.nc’, but rather than containing only output fields valid in June, it would contain all fields written between the middle of June and the middle of July, at which point one month of simulation would have elapsed, and a new output file, ‘diagnostics.2014-07.nc’, would be created.

In order to ensure that the file ‘diagnostics.2014-06.nc’ contained only data from June 2014, the reference_time attribute may be added such that the day, hour, minute, and second in the date and time represent the first day of the month at 00 UTC. In this example, the year and month of the reference time are not important, since the purpose of the reference time here is to describe to MPAS that the monthly filename interval begins (i.e., is referenced to) the first day of the month.

```
<stream name="diagnostics"
  type="output"
  filename_template="diagnostics.$Y-$M.nc"
  filename_interval="01-00_00:00:00"
  reference_time="2014-01-01_00:00:00"
  precision="single"
  clobber_mode="append"
  output_interval="6:00:00" />

  <var name="u10"/>
  <var name="v10"/>
  <var name="t2"/>
  <var name="q2"/>

</stream>
```

In general, the components of a timestamp, `YYYY-MM-DD_hh:mm:ss`, that are less significant than (i.e., to the right of) those contained in a filename template are important in a `reference_time`. For example, with a `filename_template` that contained only the year, the month component of the `reference_time` would become important to identify the month of the year on which the yearly basis for filenames would begin.

Chapter 5

Visualization

This chapter discusses visualization tools that may be used by all cores. For instructions on additional visualization tools for this core, see Chapter 8.

5.1 ParaView

ParaView may be used to visualize MPAS initialization, output, and restart files. It includes a reader that was specifically designed to read MPAS NetCDF files, including Cartesian and spherical domains. At this time, only cell-centered quantities may be plotted with ParaView. Variables located at edges and vertices must be interpolated to cell centers for visualization.

ParaView is freely available for download at <http://www.paraview.org>. Binary installations are available for Windows, Mac, and Linux, as well as source code files and tutorials. From the ParaView website:

ParaView is an open-source, multi-platform data analysis and visualization application. ParaView users can quickly build visualizations to analyze their data using qualitative and quantitative techniques. The data exploration can be done interactively in 3D or programmatically using ParaView's batch processing capabilities. ParaView was developed to analyze extremely large datasets using distributed memory computing resources. It can be run on supercomputers to analyze datasets of terascale as well as on laptops for smaller data.

To visualize an MPAS cell-centered variable in ParaView, open the file and choose **MPAS NetCDF (Unstructured)** as the file format. In the lower left Object Inspector panel, choose your variables of interest (Figure 5.1). For large data sets, loading fewer variables will result in less wait time. Options are available for latitude-longitude projections, vertical level, etc. Click the 'Apply' button to load the data set. In the toolbars at the top, choose the variable to plot from the pull-down menu, and 'Surface' for the type of visualization. The color bar button displays a color bar, and the color scale editor button allows the user to manually change the color bar type and extents. The Filters menu provides computational tools for interactive data manipulation. Movies, in avi format or as individual frames, may be conveniently created with the **Save Animation** tool in the File menu.

Paraview may be used to view the grid from any MPAS NetCDF file by choosing **Wireframe** or **Surface With Edges** from the visualization-type pull-down menu (Figure 5.2). This produces a view of the Delaunay triangulation, which is the dual mesh to the primal Voronoi cell grid (Figure

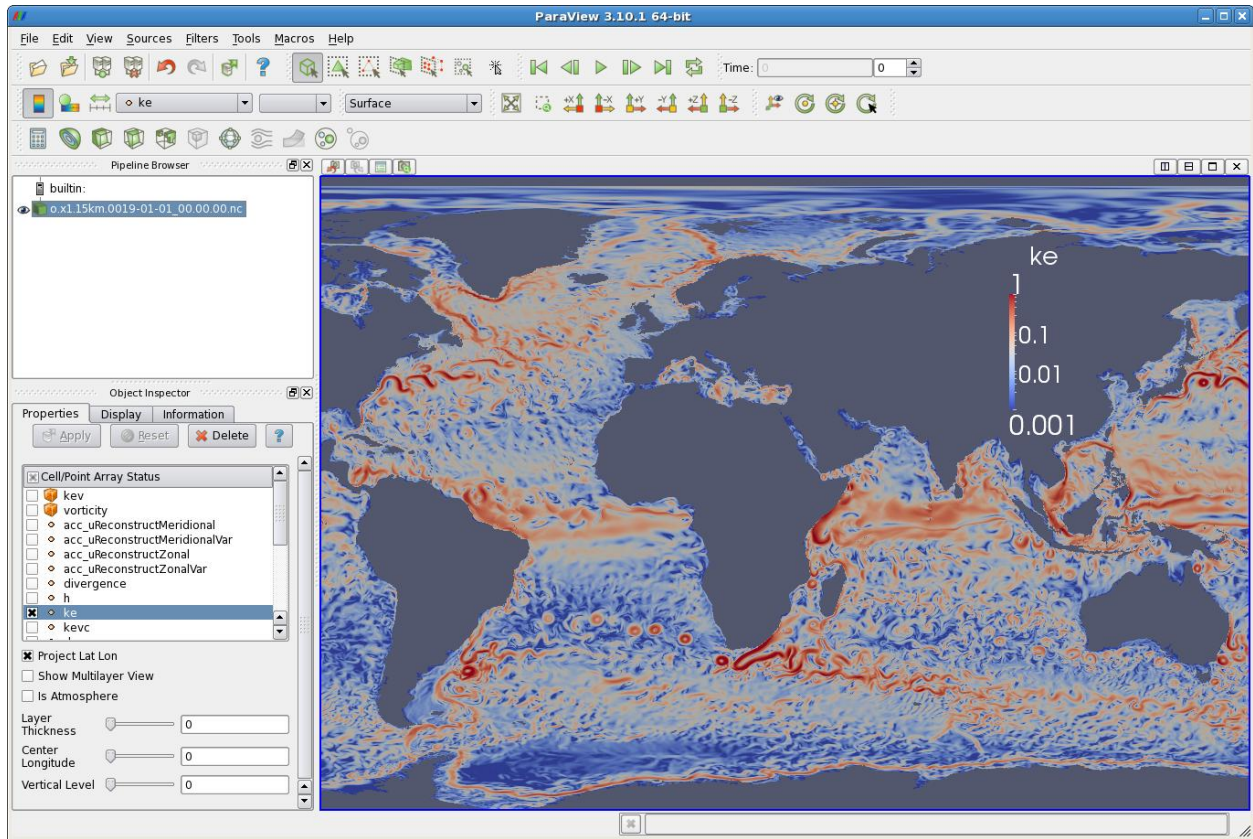


Figure 5.1: Example of ParaView to view an MPAS NetCDF file.

3.1). Paraview plots all variables by interpolating colors between each corner of the Delaunay triangles. These corners are the cell-center locations of the primal grid.




Figure 5.2: Example of visualizing the dual mesh from an MPAS NetCDF file.

Part II

MPAS-Land Ice

Chapter 6

Governing Equations

Advection is performed on a C-grid, with scalar quantities (thickness, temperature, age, etc.) on the Voronoi cell centers and velocities and fluxes centered at Voronoi cell edges. MPAS-Land Ice uses SI units everywhere, including input and output. One exception is the ability (but not the requirement) to specify the model time step in years (but which is then converted to seconds internally).

6.1 Momentum Balance

Currently within MPAS-Land Ice, the momentum balance for ice is approximated with the Shallow Ice Approximation (SIA) (Hutter, 1983), which is solved explicitly. In terms of balancing the gravitational body force, the SIA neglects all but the 0th-order, vertical shear-stress gradients. The preferred numerical approach for implementing the SIA in ice sheet models is not to solve for the velocity directly but to instead formulate a parabolic PDE describing the thickness evolution, with velocities implicit in the formulation. However, for higher-order treatments of the momentum balance, it is necessary to solve the velocity and thickness evolution steps separately. Therefore, to allow for the eventual incorporation of higher-order velocity solvers in MPAS Land Ice, the current design explicitly calculates velocities from the SIA.

Within a column, at any point in the model domain in map view, the depth-dependent SIA velocity can be solved for as:

$$\mathbf{u}(z) = -\frac{1}{2}A(\rho g)^n |\nabla s|^{n-1} \nabla s [H^{n+1} - (h - z)^{n+1}] \quad (6.1)$$

where $\mathbf{u}(z)$ is the horizontal velocity vector, A is the flow rate factor (primarily a function of ice temperature), n is the Glen flow law exponent (typically taken as 3), ρ is the density of ice, g is acceleration due to gravity, s is the ice surface elevation, H is ice thickness, and z is the vertical coordinate. Velocities are (nonlinearly) proportional to both the ice thickness and the ice surface slope.

Velocities and fluxes are calculated on the midpoint of Voronoi cell edges. The normal component of surface slope is calculated on cell edges using surface elevation at adjacent cell centers. The tangential component of surface slope is calculated on cell edges using surface elevation at adjacent vertices. The surface elevation at vertices is calculated from the values at adjacent cell centers using barycentric interpolation. Ice thickness on edges is calculated as the average of the adjacent cell center values (2nd-order approximation).

6.2 Time Integration

Currently, MPAS Land Ice only supports Forward Euler time integration.

6.3 Advection

Currently, MPAS Land Ice only supports advection of thickness and only using First-Order Upwinding. In 1D, first-order upwinding of ice thickness using a Forward Euler time step is described by:

$$\frac{H_i^{n+1} - H_i^n}{\Delta t} + u \frac{H_i^n - H_{i-1}^n}{\Delta x} = 0 \quad \text{for } u > 0 \quad (6.2)$$

$$\frac{H_i^{n+1} - H_i^n}{\Delta t} + u \frac{H_{i+1}^n - H_i^n}{\Delta x} = 0 \quad \text{for } u < 0 \quad (6.3)$$

where Δx represents the horizontal grid spacing along flow, subscripts designate the spatial dimension, and superscripts designate the time dimension.

(See, e.g. http://en.wikipedia.org/wiki/Upwind_scheme)

To allow the eventual inclusion of tracer advection, thickness is advected level-by-level, rather than the cheaper operation of advecting the total column thickness.

Chapter 7

Model Configuration

This chapter describes the configuration of the Land Ice core. The chapter covers the dimensions used in the model, the Namelist options which are used to provide run-time configurability of model options, the variables used in the model, and the usage of run-time I/O streams.

7.1 Dimensions

Name	Units	Description
nCells	<i>unitless</i>	The number of polygons in the primary grid.
nEdges	<i>unitless</i>	The number of edge midpoints in either the primary or dual grid.
maxEdges	<i>unitless</i>	The largest number of edges any polygon within the grid has.
maxEdges2	<i>unitless</i>	Two times the largest number of edges any polygon within the grid has.
nVertices	<i>unitless</i>	The total number of cells in the dual grid. Also the number of corners in the primary grid.
TWO	<i>unitless</i>	The number two as a dimension.
R3	<i>unitless</i>	The number three as a dimension.
vertexDegree	<i>unitless</i>	The number of cells or edges touching each vertex.
nVertLevels	<i>unitless</i>	The number of levels in the vertical direction. All vertical levels share the same horizontal locations.
nVertLevelsP1	<i>unitless</i>	The number of interfaces in the vertical direction.

7.2 Namelist options

Embedded links point to more detailed namelist information in the appendix.

7.2.1 velocity_solver

The velocity_solver namelist record controls which velocity solver is used and options associated with velocity solvers.

Name	Description
config_velocity_solver	Selection of the method for solving ice velocity.

7.2.2 advection

The advection namelist record controls options associated with advection of thickness and tracers. Tracer advection is not currently supported.

Name	Description
config_thickness_advection	Selection of the method for advecting thickness.
config_tracer_advection	Selection of the method for advecting tracers.

7.2.3 physical_parameters

The physical_parameters namelist record sets scalar physical parameters and constants within the land ice model.

Name	Description
config_ice_density	ice density to use
config_ocean_density	ocean density to use for calculating floatation
config_sea_level	sea level to use for calculating floatation
config_default_flowParamA	Defines the default value of the flow law parameter A to be used if it is not being calculated from ice temperature. Defaults to the SI representation of $1.0\text{e-}16 \text{ yr}^{-1} \text{ Pa}^{-3}$.
config_flowLawExponent	Defines the value of the Glen flow law exponent, n.
config_dynamic_thickness	Defines the ice thickness below which dynamics are not calculated.

7.2.4 time_integration

The time integration namelist record controls parameters that pertain to all time-stepping methods. At present, Forward Euler is the only time integration method implemented.

Name	Description
config_dt	Length of model time step defined as a time interval.
config_time_integration	Time integration method.

7.2.5 time_management

General time management is handled by the time_management namelist record. Included options handle time-related parts of MPAS, such as the calendar type and if the simulation is a restart or not.

Users should use this record to specify the beginning time of the simulation, and either the duration or the end of the simulation. Only the end or the duration need to be specified as the other is derived within MPAS from the beginning time and other specified one.

If both the run duration and stop time are specified, run duration is used in place of stop time.

Name	Description
config_do_restart	Determines if the initial conditions should be read from a restart file, or an input file. To perform a restart, simply set this to true in the namelist.input file and modify the start time to be the time you want restart from. A restart will read the grid information from the input field, and the restart state from the restart file. It will perform a run normally, except velocity will not be solved on a restart.
config_restart_timestamp_name	Path to the filename for restart timestamps to be read and written from.
config_start_time	Timestamp describing the initial time of the simulation. If it is set to 'file', the initial time is read from restart_timestamp
config_stop_time	Timestamp describing the final time of the simulation. If it is set to 'none' the final time is determined from config_start_time and config_run_duration. If config_run_duration is also specified, it takes precedence over config_stop_time. Set config_stop_time to be equal to config_start_time (and config_run_duration to 'none') to perform a diagnostic solve only.
config_run_duration	Timestamp describing the length of the simulation. If it is set to 'none' the duration is determined from config_start_time and config_stop_time. config_run_duration overrides inconsistent values of config_stop_time. If a time value is specified for config_run_duration, it must be greater than 0.
config_calendar_type	Selection of the type of calendar that should be used in the simulation.

7.2.6 io

The io namelist record provides options for modifications to the I/O system of MPAS. These include frequency, file name, and parallelization options.

Name	Description
config_write_output_on_startu-p	Logical flag determining if an output file should be written prior to the first time step.
config_pio_num_iotasks	Integer specifying how many IO tasks should be used within the PIO library. A value of 0 causes all MPI tasks to also be IO tasks. IO tasks are required to write contiguous blocks of data to a file.
config_pio_stride	Integer specifying the stride of each IO task.
config_year_digits	Integer specifying the number of digits used to represent the year in time strings.

7.2.7 decomposition

MPAS handles decomposing all variables into computational blocks. The decomposition used needs to be specified at run time and is computed by an external tool (e.g. metis). Additionally, MPAS supports multiple computational blocks per MPI process, and the user may specify an additional decomposition file which can specify the assignment of blocks to MPI processes. Run-time parameters that control the run-time decomposition used are specified within the decomposition namelist record.

Name	Description
config_num_halos	Determines the number of halo cells extending from a blocks owned cells (Called the 0-Halo). The default of 3 is the minimum that can be used with monotonic advection.
config_block_decomp_file_prefix	Defines the prefix for the block decomposition file. Can include a path. The number of blocks is appended to the end of the prefix at run-time.
config_number_of_blocks	Determines the number of blocks a simulation should be run with. If it is set to 0, the number of blocks is the same as the number of MPI tasks at run-time.
config_explicit_proc_decomp	Determines if an explicit processor decomposition should be used. This is only useful if multiple blocks per processor are used.
config_proc_decomp_file_prefix	Defines the prefix for the processor decomposition file. This file is only read if config_explicit_proc_decomp is .true. The number of processors is appended to the end of the prefix at run-time.

7.2.8 debug

At run-time a user can enable debugging features within MPAS-Land Ice. Currently the only debug option is to print more detailed information about thickness advection. Potential future debug options would be to include disabling of any tendencies to help determine why an issue might be happening; various checks on certain fields; and the ability to prescribe both a thickness and velocity field at run-time which are constant throughout a simulation. All options that control these debugging features are specified within the debug namelist record.

Name	Description
config_print_thickness_advection_info	Prints additional information about thickness advection.

7.3 Variable definitions

Embedded links point to more detailed variable information in the appendix.

7.3.1 [state](#)

Name	Description
xtime	model time, with format 'YYYY-MM-DD_HH:MM:SS'
thickness	ice thickness
layerThickness	layer thickness
temperature	ice temperature
lowerSurface	elevation at bottom of ice
upperSurface	elevation at top of ice
layerThicknessEdge	layer thickness on cell edges
upperSurfaceVertex	elevation at top of ice on vertices (currently only needed by shallow ice solver)
cellMask	bitmask indicating various properties about the ice sheet on cells. cellMask only needs to be a restart field if <code>config.allow_additional_advance = false</code> (to keep the mask of initial ice extent)
edgeMask	bitmask indicating various properties about the ice sheet on edges.
vertexMask	bitmask indicating various properties about the ice sheet on vertices.
normalVelocity	horizontal velocity, normal component to an edge
uReconstructX	x-component of velocity reconstructed on cell centers
uReconstructY	y-component of velocity reconstructed on cell centers
uReconstructZ	z-component of velocity reconstructed on cell centers
uReconstructZonal	zonal velocity reconstructed on cell centers
uReconstructMeridional	meridional velocity reconstructed on cell centers

7.3.2 [tend](#)

Name	Description
tend_layerThickness	time tendency of layer thickness
tend_temperature	time tendency of ice temperature

7.3.3 mesh

Name	Description
latCell	Latitude location of cell centers in radians.
lonCell	Longitude location of cell centers in radians.
xCell	X Coordinate in cartesian space of cell centers.
yCell	Y Coordinate in cartesian space of cell centers.
zCell	Z Coordinate in cartesian space of cell centers.
indexToCellID	List of global cell IDs.
latEdge	Latitude location of edge midpoints in radians.
lonEdge	Longitude location of edge midpoints in radians.
xEdge	X Coordinate in cartesian space of edge midpoints.
yEdge	Y Coordinate in cartesian space of edge midpoints.
zEdge	Z Coordinate in cartesian space of edge midpoints.
indexToEdgeID	List of global edge IDs.
latVertex	Latitude location of vertices in radians.
lonVertex	Longitude location of vertices in radians.
xVertex	X Coordinate in cartesian space of vertices.
yVertex	Y Coordinate in cartesian space of vertices.
zVertex	Z Coordinate in cartesian space of vertices.
indexToVertexID	List of global vertex IDs.
cellsOnEdge	List of cells that straddle each edge.
nEdgesOnCell	Number of edges that border each cell.
nEdgesOnEdge	Number of edges that surround each of the cells that straddle each edge. These edges are used to reconstruct the tangential velocities.
edgesOnCell	List of edges that border each cell.
edgesOnEdge	List of edges that border each of the cells that straddle each edge.
weightsOnEdge	Reconstruction weights associated with each of the edgesOnEdge.
dvEdge	Length of each edge, computed as the distance between verticesOnEdge.
dcEdge	Length of each edge, computed as the distance between cellsOnEdge.
angleEdge	Angle the edge normal makes with local eastward direction.
areaCell	Area of each cell in the primary grid.
areaTriangle	Area of each cell (triangle) in the dual grid.
edgeNormalVectors	Normal vector defined at an edge.
localVerticalUnitVectors	Unit surface normal vectors defined at cell centers.
cellTangentPlane	The two vectors that define a tangent plane at a cell center.
cellsOnCell	List of cells that neighbor each cell.
verticesOnCell	List of vertices that border each cell.
verticesOnEdge	List of vertices that straddle each edge.
edgesOnVertex	List of edges that share a vertex as an endpoint.
cellsOnVertex	List of cells that share a vertex.
kiteAreasOnVertex	Area of the portions of each dual cell that are part of each cellsOnVertex.
coeffs_reconstruct	Coefficients to reconstruct velocity vectors at cells centers.
edgeSignOnCell	Sign of edge contributions to a cell for each edge on cell. Used for bit-reproducible loops. Represents directionality of vector connecting cells.

Name	Description (Continued)
edgeSignOnVertex	Sign of edge contributions to a vertex for each edge on vertex. Used for bit-reproducible loops. Represents directionality of vector connecting vertices.
layerThicknessFractions	Fractional thickness of each sigma layer
layerCenterSigma	Sigma (fractional) level at center of each layer
layerInterfaceSigma	Sigma (fractional) level at interface between each layer (including top and bottom)
bedTopography	Elevation of ice sheet bed. Once isostasy is added to the model, this should become a state variable.
sfcMassBal	Surface mass balance

7.4 Run-time input/output streams

Chapter 4 provides a detailed overview of the implementation of run-time input/output streams in MPAS. Within the Land Ice core, the following streams are defined at build time:

7.4.1 input

This is an immutable stream defining the fields required for input. It is only read at the initial time. Input files may have other fields in them, but only the fields specified in this stream definition are actually read. Default name is `landice_grid.nc`.

The input stream consists of the following members:

- stream name="basicmesh"
- var_array name="tracers"
- var name="thickness"
- var name="normalVelocity"
- var name="bedTopography"
- var name="sfcMassBal"

7.4.2 output

This is a mutable stream defining the fields that will be output. Because it is mutable, the list of fields for output may be modified at run time by editing the `streams.landice` file. Default name is `output.nc`. Default clobber mode is `replace_files`, **which will overwrite existing output**.

The output stream consists of the following members by default:

- stream name="basicmesh"
- var_array name="tracers"
- var name="xtime"
- var name="thickness"

- var name="layerThickness"
- var name="lowerSurface"
- var name="upperSurface"
- var name="cellMask"
- var name="edgeMask"
- var name="vertexMask"
- var name="normalVelocity"
- var name="uReconstructX"
- var name="uReconstructY"
- var name="uReconstructZ"
- var name="uReconstructZonal"
- var name="uReconstructMeridional"

7.4.3 restart

This is an immutable stream defining the fields required for restart. It is both an input and output stream. The model writes restart files with a single time level in them periodically. If a restart from one of these checkpoints is desired, set `config_do_restart` to `.true.` and set `config_start_time` to `file` in `namelist.landice`. The model will take the start time from the value in the text file specified by `config_restart_timestamp_name` (default name is "restart_timestamp"), and use the associated `.nc` file to restart the model from that checkpoint.

Default name is `restart.$Y-$M-$D_$h.$m.$s.nc` with a new file for each checkpoint. Default clobber mode is `replace_files`, **which will overwrite existing output**.

The user does not need to keep track of what fields are required for restart, but for reference they are:

- stream name="basicmesh"
- var_array name="tracers"
- var name="xtime"
- var name="thickness"
- var name="cellMask"
- var name="normalVelocity"
- var name="bedTopography"
- var name="sfcMassBal"

7.4.4 basicmesh

This is an immutable stream that specifies the list of fields that make up the MPAS mesh specification. It is provided as a convenience for including mesh fields in other streams without having to list them all explicitly.

7.4.5 Other streams

As described in Chapter 4, additional streams may be added by the user at run-time. One common example would be a “forcing” stream that gets read at each time level while the model runs. This can be accomplished by creating an input stream with `input_interval` set to a time interval. If the model does not find the current time in the forcing file, it will read the latest value before the current time instead (piecewise constant forcing).

Chapter 8

Land Ice Visualization

This chapter discusses visualization tools that are specific to the Land Ice core. For instructions on visualization tools that may be used by all cores, such as Paraview, see Chapter 5.

8.1 Python

Python visualization scripts are available for the dome test case, and general python visualization tools are in development. In order to use these scripts, the following python modules are required:

- matplotlib, see <http://matplotlib.org>
- numpy, see <http://www.numpy.org>
- pylab, see www.scipy.org
- netCDF4, see <http://code.google.com/p/netcdf4-python>

Most package managers (including MacPorts) have packages for these python modules. Another convenient way to install all these libraries at once is to purchase the Enthought Python Distribution (EPD), available at <https://www.enthought.com/products/epd>. Many institutions have Python-EPD installed on their compute clusters.

Chapter 9

Test Cases

Eventually test cases will be available for download. Currently they are only part of the Development code for MPAS-Land Ice.

9.1 Halfar Dome

This test case describes the time evolution of a dome of ice as described by Halfar (1983). This test provide an analytic solution for a flat-bedded SIA problem.

$$\frac{\partial H}{\partial t} = \nabla \cdot (\Gamma H^{n+2} |\nabla H|^{n-1} \nabla H) \quad (9.1)$$

where n is the exponent in the Glen flow law, commonly taken as 3, and Γ is a positive constant:

$$\Gamma = \frac{2}{n+2} A(\rho g)^n \quad (9.2)$$

For $n = 3$, this reduces to:

$$H(t, r) = H_0 \left(\frac{t_0}{t} \right)^{\frac{1}{9}} \left[1 - \left(\left(\frac{t_0}{t} \right)^{\frac{1}{18}} \frac{r}{R_0} \right)^{\frac{4}{3}} \right]^{\frac{3}{7}} \quad (9.3)$$

where

$$t_0 = \frac{1}{18\Gamma} \left(\frac{7}{4} \right)^3 \frac{R_0^4}{H_0^7} \quad (9.4)$$

and H_0, R_0 are the central height of the dome and its radius at time $t = t_0$.

For more details see <http://www.projects.science.uu.nl/iceclimate/karthauss/2009/more/lecturenotes/EdBueler.pdf>, Bueler et al. (2005), Halfar (1983).

9.1.1 Provided Files

Our implementation of the Halfar dome has an initial radius of $R_0 = 21.2$ km and an initial thickness of $H = 707.1$ m. These values can be changed by editing `setup_dome_initial_conditions.py`.

- README:
Information about the test case.

- `namelist.landice`:
This file is used for actually running the dome test case in the MPAS land ice core. It may not include all options available to the model. See the `namelist.landice.defaults` file in the MPAS root directory for a list of all options available. They are also documented in [Section 7.2](#).
- `streams.landice`:
This file is used for specifying file input/output settings for the model.
- `halfar.py`:
This is the script to compare model results to the analytic solution.
- `visualize_dome.py`:
This python script provides some general visualization of the model output. It can be used in addition to `halfar.py` for additional visualization.
- `namelist.input.periodic.hex`:
This file is used for running the grid generation tool `periodic_hex` to create a grid for the test case. It should be renamed to `namelist.config` when executing `periodic_hex`. `periodic_hex` will be used to generate `grid.nc` (which can be used to create `landice_grid.nc`) and `graph.info.part.*` which can be used for running the model on more than one processor.
- `setup_dome_initial_conditions.py`:
This python script generates the dome initial condition after an empty `landice_grid.nc` file exists. If you downloaded a tar archive, you do not need to do this. However, if you want to modify the IC for some reason, you can edit and run this script.

9.1.2 Results

As the dome of ice evolves, its margin advances and its thickness decreases (there is no surface mass balance to add new mass). The script `halfar.py` will plot the modeled and analytic thickness at a specified time ([Figure 9.1](#)), as well as report model error statistics. Invoke `halfar.py --help` for details of its usage.




Figure 9.1: Halfar test case results after 200 years of dome evolution. This figure is generated by `halfar.py`.

9.2 EISMINT-1 Test Cases

This test case is from the European Ice Sheet Modelling INiTiative intercomparison experiments. These experiments are described at <http://homepages.vub.ac.be/~phuybrec/eismint.html> and in Huybrechts et al. (1996).

Currently only the Moving Margin 1 Test Case from EISMINT-1 is included.

9.2.1 Provided Files

- `namelist.landice`:
This file is used for actually running the dome test case in the MPAS land ice core. It may not include all options available to the model. See the `namelist.landice.defaults` file in the MPAS root directory for a list of all options available. They are also documented in Section 7.2.
- `streams.landice`:
This file is used for specifying file input/output settings for the model.
- `check_output_eismint-mm1.py`
This script can be used to compare model output to results from the EISMINT intercomparison.
- `namelist.input.periodic.hex`
This file is used for running the grid generation tool `periodic.hex` to create a grid for the test case. It needs to be renamed to 'namelist.input' to run `periodic.hex` if mesh needs to be generated. If you downloaded a tar archive of this test case, you do not need to create the mesh and can ignore this file.
- `setup_initial_conditions_EISMINT1-MovingMargin-1.py`
This file can be used to setup the initial conditions for the test case. If you downloaded a tar archive, you do not need to do this. However, if you want to modify the IC for some reason, you can edit and run this script.

9.2.2 Results

As the initial ice sheet evolves, its shape eventually reaches a steady-state with the imposed surface mass balance. The script `check_output_eismint-mm1.py` will plot the modeled thickness at a specified time, as well as compare the model results to the results from the original EISMINT intercomparison. Invoke `check_output_eismint-mm1.py --help` for details of its usage. The script will compare the maximum ice thickness at the final time of the model output to the values reported from the models participating in the EISMINT-1 intercomparison. You should see something similar to this:

```
=====
Max modeled thickness (m) = 2974.79474126
EISMINT models ice thickness at divide (m):
  3d models (10 of them): 2978.0 +/- 19.3
  2d models (3 of them): 2982.2 +/- 26.4
=====
```

9.3 Real World Test Cases

Eventually grids for real-world Greenland and Antarctica will be provided at varying resolutions.

Chapter 10

Global Statistics

Eventually global statistics will be calculated within MPAS Land Ice.

Chapter 11

Running MPAS-Land Ice within a coupled climate model

Eventually MPAS-Land Ice will be coupled within global climate models.

Chapter 12

Troubleshooting

12.1 Choice of time step

Symptoms: “Error in calculating thickness tendency (possibly CFL violation)” appears in log.0000.err file.

Possible cause: Time step is too long.

Remedy: Shorten time step.

Discussion: The time step must be short enough that the CFL criterion is not violated. Eventually an adaptive time integrator will be added to MPAS-Land Ice.

Chapter 13

Known Issues

- The barycentric interpolation used to calculate surface elevation at vertices gives garbage values for vertices associated with obtuse triangles on the dual mesh. Therefore, the model will only work properly for meshes with no obtuse triangles. Currently there is no error message when this occurs, so users must be aware of this constraint. Future work will improve the barycentric interpolation method to work for obtuse triangles.
- Paraview plots periodic fields in a messy way with lines connecting the periodic cells across the domain.
- Paraview gives the following fatal error with some Land Ice output files: "NetCDF: Start+count exceeds dimension bound".
- Paraview will not recognize fields without a vertical dimension (e.g. thickness will not be recognized) in versions earlier than 4.1.

Part III

Bibliography

Bibliography

- Bueler, E., C. S. Lingle, J. a. Kallen-Brown, D. N. Covey, and L. N. Bowman, 2005: Exact solutions and verification of numerical models for isothermal ice sheets. *Journal of Glaciology*, **51**, 291–306, doi:10.3189/172756505781829449.
URL <http://openurl.ingenta.com/content/xref?genre=article&issn=0022-1430&volume=51&issue=173&spage=291>
- Edwards, T. L., X. Fettweis, O. Gagliardini, F. Gillet-Chaulet, H. Goelzer, J. M. Gregory, M. Hoffman, P. Huybrechts, A. J. Payne, M. Perego, S. Price, A. Quiquet, and C. Ritz., 2013: Effect of uncertainty in surface mass balance elevation feedback on projections of the future sea level contribution of the Greenland ice sheet - Part 2: Projections. *The Cryosphere Discussions*, **7**, 675–708.
- Halfar, P., 1983: On the Dynamics of the Ice Sheets 2. *Journal of Geophysical Research*, **88**, 6043–6051.
- Hutter, K., 1983: *Theoretical glaciology; material science of ice and the mechanics of glaciers and ice sheets*. Reidel Publishing Co., Terra Scientific Publishing Co., Tokyo.
- Huybrechts, P., T. Payne, and T. E. I. Group, 1996: The EISMINT benchmarks for testing ice-sheet models. *Annals of Glaciology*, **23**, 1–12.
- Leng, W., L. Ju, M. Gunzburger, S. Price, and T. Ringler, 2012: A parallel high-order accurate finite element nonlinear Stokes ice sheet model and benchmark experiments. *Journal of Geophysical Research*, **117**, F01001, doi:10.1029/2011JF001962.
- Perego, M., M. Gunzburger, and J. Burkardt, 2012: Parallel finite-element implementation for higher-order ice-sheet models. *Journal of Glaciology*, **58**, 76–88, doi:10.3189/2012JoG11J063.
URL <http://www.igsoc.org/journal/current/207/t11J063.pdf>
- Shannon, S. R., A. J. Payne, I. D. Bartholomew, M. R. V. D. Broeke, T. L. Edwards, A. J. Sole, R. S. W. V. D. Wal, and T. Zwinger, 2013: Enhanced basal lubrication and the contribution of the Greenland ice sheet to future sea-level rise. *Proceedings of the National Academy of Sciences of the United States of America*, 1–6, doi:10.1073/pnas.1212647110.

Part IV

Appendices

Appendix A

Namelist options

A.1 config_block_decomp_file_prefix

Type:	character
Units:	<i>unitless</i>
Default Value:	graph.info.part.
Possible Values:	Any path/prefix to a block decomposition file.

Table A.1: config_block_decomp_file_prefix: Defines the prefix for the block decomposition file. Can include a path. The number of blocks is appended to the end of the prefix at run-time.

A.2 config_calendar_type

Type:	character
Units:	<i>unitless</i>
Default Value:	gregorian_noleap
Possible Values:	'gregorian', 'gregorian_noleap'

Table A.2: config_calendar_type: Selection of the type of calendar that should be used in the simulation.

A.3 config_default_flowParamA

Type:	real
Units:	$s^{-1} Pa^{-n}$
Default Value:	3.1709792e-24
Possible Values:	Any positive real value

Table A.3: config_default_flowParamA: Defines the default value of the flow law parameter A to be used if it is not being calculated from ice temperature. Defaults to the SI representation of $1.0\text{e-}16 \text{ yr}^{-1} \text{ Pa}^{-3}$.

A.4 config_do_restart

Type:	logical
Units:	<i>unitless</i>
Default Value:	.false.
Possible Values:	.true. or .false.

Table A.4: config_do_restart: Determines if the initial conditions should be read from a restart file, or an input file. To perform a restart, simply set this to true in the namelist.input file and modify the start time to be the time you want restart from. A restart will read the grid information from the input field, and the restart state from the restart file. It will perform a run normally, except velocity will not be solved on a restart.

A.5 config_dt

Type:	character
Units:	<i>unitless</i>
Default Value:	0001-00-00_00:00:00
Possible Values:	Any time interval of the format 'YYYY-MM-DD.HH:MM:SS', but limited by CFL condition. (items in the format string may be dropped from the left if not needed, and the components on either side of the underscore may be replaced with a single integer representing the rightmost unit)

Table A.5: config_dt: Length of model time step defined as a time interval.

A.6 config_dynamic_thickness

Type:	real
Units:	<i>m of ice</i>
Default Value:	100.0

Possible Values:	Any positive real value
------------------	-------------------------

Table A.6: `config_dynamic_thickness`: Defines the ice thickness below which dynamics are not calculated.

A.7 `config_explicit_proc_decomp`

Type:	logical
Units:	<i>unitless</i>
Default Value:	<code>.false.</code>
Possible Values:	<code>.true.</code> or <code>.false.</code>

Table A.7: `config_explicit_proc_decomp`: Determines if an explicit processor decomposition should be used. This is only useful if multiple blocks per processor are used.

A.8 `config_flowLawExponent`

Type:	real
Units:	<i>none</i>
Default Value:	3.0
Possible Values:	Any real value

Table A.8: `config_flowLawExponent`: Defines the value of the Glen flow law exponent, n .

A.9 `config_ice_density`

Type:	real
Units:	$kg\ m^{-3}$
Default Value:	910.0
Possible Values:	Any positive real value

Table A.9: `config_ice_density`: ice density to use

A.10 `config_num_halos`

Type:	integer
Units:	<i>unitless</i>
Default Value:	3
Possible Values:	Any positive interger value.

Table A.10: `config_num_halos`: Determines the number of halo cells extending from a blocks owned cells (Called the 0-Halo). The default of 3 is the minimum that can be used with monotonic advection.

A.11 `config_number_of_blocks`

Type:	integer
Units:	<i>unitless</i>
Default Value:	0
Possible Values:	Any integer ≥ 0 .

Table A.11: `config_number_of_blocks`: Determines the number of blocks a simulation should be run with. If it is set to 0, the number of blocks is the same as the number of MPI tasks at run-time.

A.12 `config_ocean_density`

Type:	real
Units:	$kg\ m^{-3}$
Default Value:	1028.0
Possible Values:	Any positive real value

Table A.12: `config_ocean_density`: ocean density to use for calculating floatation

A.13 `config_pio_num_iotasks`

Type:	integer
Units:	<i>unitless</i>
Default Value:	0
Possible Values:	Any positive integer value greater than or equal to 0.

Table A.13: `config_pio_num_iotasks`: Integer specifying how many IO tasks should be used within the PIO library. A value of 0 causes all MPI tasks to also be IO tasks. IO tasks are required to write contiguous blocks of data to a file.

A.14 config_pio_stride

Type:	integer
Units:	<i>unitless</i>
Default Value:	1
Possible Values:	Any positive integer value greater than 0.

Table A.14: config_pio_stride: Integer specifying the stride of each IO task.

A.15 config_print_thickness_advection_info

Type:	logical
Units:	<i>unitless</i>
Default Value:	.false.
Possible Values:	.true. or .false.

Table A.15: config_print_thickness_advection_info: Prints additional information about thickness advection.

A.16 config_proc_decomp_file_prefix

Type:	character
Units:	<i>unitless</i>
Default Value:	graph.info.part.
Possible Values:	Any path/prefix to a processor decomposition file.

Table A.16: config_proc_decomp_file_prefix: Defines the prefix for the processor decomposition file. This file is only read if config_explicit_proc_decomp is .true. The number of processors is appended to the end of the prefix at run-time.

A.17 config_restart_timestamp_name

Type:	character
Units:	<i>unitless</i>
Default Value:	restart_timestamp

Possible Values:	Path to a file.
------------------	-----------------

Table A.17: `config_restart_timestamp_name`: Path to the filename for restart timestamps to be read and written from.

A.18 `config_run_duration`

Type:	character
Units:	<i>unitless</i>
Default Value:	none
Possible Values:	'YYYY-MM-DD_HH:MM:SS' or 'none' (items in the format string may be dropped from the left if not needed, and the components on either side of the underscore may be replaced with a single integer representing the rightmost unit)

Table A.18: `config_run_duration`: Timestamp describing the length of the simulation. If it is set to 'none' the duration is determined from `config_start_time` and `config_stop_time`. `config_run_duration` overrides inconsistent values of `config_stop_time`. If a time value is specified for `config_run_duration`, it must be greater than 0.

A.19 `config_sea_level`

Type:	real
Units:	<i>m above datum</i>
Default Value:	0.0
Possible Values:	Any real value

Table A.19: `config_sea_level`: sea level to use for calculating floatation

A.20 `config_start_time`

Type:	character
Units:	<i>unitless</i>
Default Value:	0000-01-01_00:00:00

Possible Values:	'YYYY-MM-DD_HH:MM:SS' (items in the format string may be dropped from the left if not needed, and the components on either side of the underscore may be replaced with a single integer representing the rightmost unit)
------------------	--

Table A.20: `config_start_time`: Timestamp describing the initial time of the simulation. If it is set to 'file', the initial time is read from `restart_timestamp`

A.21 `config_stop_time`

Type:	character
Units:	<i>unitless</i>
Default Value:	0000-01-01_00:00:00
Possible Values:	'YYYY-MM-DD_HH:MM:SS' or 'none' (items in the format string may be dropped from the left if not needed, and the components on either side of the underscore may be replaced with a single integer representing the rightmost unit)

Table A.21: `config_stop_time`: Timestamp describing the final time of the simulation. If it is set to 'none' the final time is determined from `config_start_time` and `config_run_duration`. If `config_run_duration` is also specified, it takes precedence over `config_stop_time`. Set `config_stop_time` to be equal to `config_start_time` (and `config_run_duration` to 'none') to perform a diagnostic solve only.

A.22 `config_thickness_advection`

Type:	character
Units:	<i>unitless</i>
Default Value:	fo
Possible Values:	'fo', 'none'

Table A.22: `config_thickness_advection`: Selection of the method for advecting thickness.

A.23 `config_time_integration`

Type:	character
Units:	<i>unitless</i>
Default Value:	forward_euler

Possible Values:	'forward_euler'
------------------	-----------------

Table A.23: config_time_integration: Time integration method.

A.24 config_tracer_advection

Type:	character
Units:	<i>unitless</i>
Default Value:	none
Possible Values:	'none'

Table A.24: config_tracer_advection: Selection of the method for advecting tracers.

A.25 config_velocity_solver

Type:	character
Units:	<i>unitless</i>
Default Value:	sia
Possible Values:	'sia'

Table A.25: config_velocity_solver: Selection of the method for solving ice velocity.

A.26 config_write_output_on_startup

Type:	logical
Units:	<i>unitless</i>
Default Value:	.true.
Possible Values:	.true. or .false.

Table A.26: config_write_output_on_startup: Logical flag determining if an output file should be written prior to the first time step.

A.27 config_year_digits

Type:	integer
-------	---------

Units:	<i>unitless</i>
Default Value:	4
Possible Values:	Any positive integer value greater than 0.

Table A.27: `config_year_digits`: Integer specifying the number of digits used to represent the year in time strings.

Appendix B

Variable definitions

B.1 mesh

B.1.1 angleEdge

Type:	real
Units:	<i>radians</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'angleEdge' in 'mesh' pool

Table B.1: angleEdge: Angle the edge normal makes with local eastward direction.

B.1.2 areaCell

Type:	real
Units:	m^2
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'areaCell' in 'mesh' pool

Table B.2: areaCell: Area of each cell in the primary grid.

B.1.3 areaTriangle

Type:	real
-------	------

Units:	m^2
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'areaTriangle' in 'mesh' pool

Table B.3: areaTriangle: Area of each cell (triangle) in the dual grid.

B.1.4 bedTopography

Type:	real
Units:	<i>m above datum</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	input restart
Pool path:	'bedTopography' in 'mesh' pool

Table B.4: bedTopography: Elevation of ice sheet bed. Once isostasy is added to the model, this should become a state variable.

B.1.5 cellTangentPlane

Type:	real
Units:	<i>unitless</i>
Dimension:	R3 TWO nCells
Persistence:	persistent
Number of time levels:	1
Pool path:	'cellTangentPlane' in 'mesh' pool

Table B.5: cellTangentPlane: The two vectors that define a tangent plane at a cell center.

B.1.6 cellsOnCell

Type:	integer
Units:	<i>unitless</i>
Dimension:	maxEdges nCells
Persistence:	persistent
Number of time levels:	1

In streams:	basicmesh
Pool path:	'cellsOnCell' in 'mesh' pool

Table B.6: cellsOnCell: List of cells that neighbor each cell.

B.1.7 cellsOnEdge

Type:	integer
Units:	<i>unitless</i>
Dimension:	TWO nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'cellsOnEdge' in 'mesh' pool

Table B.7: cellsOnEdge: List of cells that straddle each edge.

B.1.8 cellsOnVertex

Type:	integer
Units:	<i>unitless</i>
Dimension:	vertexDegree nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'cellsOnVertex' in 'mesh' pool

Table B.8: cellsOnVertex: List of cells that share a vertex.

B.1.9 coeffs_reconstruct

Type:	real
Units:	<i>unitless</i>
Dimension:	R3 maxEdges nCells
Persistence:	persistent
Number of time levels:	1
Pool path:	'coeffs_reconstruct' in 'mesh' pool

Table B.9: coeffs_reconstruct: Coefficients to reconstruct velocity vectors at cells centers.

B.1.10 dcEdge

Type:	real
Units:	m
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'dcEdge' in 'mesh' pool

Table B.10: dcEdge: Length of each edge, computed as the distance between cellsOnEdge.

B.1.11 dvEdge

Type:	real
Units:	m
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'dvEdge' in 'mesh' pool

Table B.11: dvEdge: Length of each edge, computed as the distance between verticesOnEdge.

B.1.12 edgeNormalVectors

Type:	real
Units:	<i>unitless</i>
Dimension:	R3 nEdges
Persistence:	persistent
Number of time levels:	1
Pool path:	'edgeNormalVectors' in 'mesh' pool

Table B.12: edgeNormalVectors: Normal vector defined at an edge.

B.1.13 edgeSignOnCell

Type:	integer
Units:	<i>unitless</i>
Dimension:	maxEdges nCells
Persistence:	persistent
Number of time levels:	1
Pool path:	'edgeSignOnCell' in 'mesh' pool

Table B.13: edgeSignOnCell: Sign of edge contributions to a cell for each edge on cell. Used for bit-reproducible loops. Represents directionality of vector connecting cells.

B.1.14 edgeSignOnVertex

Type:	integer
Units:	<i>unitless</i>
Dimension:	maxEdges nVertices
Persistence:	persistent
Number of time levels:	1
Pool path:	'edgeSignOnVertex' in 'mesh' pool

Table B.14: edgeSignOnVertex: Sign of edge contributions to a vertex for each edge on vertex. Used for bit-reproducible loops. Represents directionality of vector connecting vertices.

B.1.15 edgesOnCell

Type:	integer
Units:	<i>unitless</i>
Dimension:	maxEdges nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'edgesOnCell' in 'mesh' pool

Table B.15: edgesOnCell: List of edges that border each cell.

B.1.16 edgesOnEdge

Type:	integer
Units:	<i>unitless</i>

Dimension:	maxEdges2 nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'edgesOnEdge' in 'mesh' pool

Table B.16: edgesOnEdge: List of edges that border each of the cells that straddle each edge.

B.1.17 edgesOnVertex

Type:	integer
Units:	<i>unitless</i>
Dimension:	vertexDegree nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'edgesOnVertex' in 'mesh' pool

Table B.17: edgesOnVertex: List of edges that share a vertex as an endpoint.

B.1.18 indexToCellID

Type:	integer
Units:	<i>unitless</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'indexToCellID' in 'mesh' pool

Table B.18: indexToCellID: List of global cell IDs.

B.1.19 indexToEdgeID

Type:	integer
Units:	<i>unitless</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1

In streams:	basicmesh
Pool path:	'indexToEdgeID' in 'mesh' pool

Table B.19: indexToEdgeID: List of global edge IDs.

B.1.20 indexToVertexID

Type:	integer
Units:	<i>unitless</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'indexToVertexID' in 'mesh' pool

Table B.20: indexToVertexID: List of global vertex IDs.

B.1.21 kiteAreasOnVertex

Type:	real
Units:	m^2
Dimension:	vertexDegree nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'kiteAreasOnVertex' in 'mesh' pool

Table B.21: kiteAreasOnVertex: Area of the portions of each dual cell that are part of each cellsOnVertex.

B.1.22 latCell

Type:	real
Units:	<i>radians</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'latCell' in 'mesh' pool

Table B.22: latCell: Latitude location of cell centers in radians.

B.1.23 latEdge

Type:	real
Units:	<i>radians</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'latEdge' in 'mesh' pool

Table B.23: latEdge: Latitude location of edge midpoints in radians.

B.1.24 latVertex

Type:	real
Units:	<i>radians</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'latVertex' in 'mesh' pool

Table B.24: latVertex: Latitude location of vertices in radians.

B.1.25 layerCenterSigma

Type:	real
Units:	<i>none</i>
Dimension:	nVertLevels
Persistence:	persistent
Number of time levels:	1
Pool path:	'layerCenterSigma' in 'mesh' pool

Table B.25: layerCenterSigma: Sigma (fractional) level at center of each layer

B.1.26 layerInterfaceSigma

Type:	real
-------	------

Units:	<i>none</i>
Dimension:	nVertLevelsP1
Persistence:	persistent
Number of time levels:	1
Pool path:	'layerInterfaceSigma' in 'mesh' pool

Table B.26: layerInterfaceSigma: Sigma (fractional) level at interface between each layer (including top and bottom)

B.1.27 layerThicknessFractions

Type:	real
Units:	<i>none</i>
Dimension:	nVertLevels
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'layerThicknessFractions' in 'mesh' pool

Table B.27: layerThicknessFractions: Fractional thickness of each sigma layer

B.1.28 localVerticalUnitVectors

Type:	real
Units:	<i>unitless</i>
Dimension:	R3 nCells
Persistence:	persistent
Number of time levels:	1
Pool path:	'localVerticalUnitVectors' in 'mesh' pool

Table B.28: localVerticalUnitVectors: Unit surface normal vectors defined at cell centers.

B.1.29 lonCell

Type:	real
Units:	<i>radians</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh

Pool path:	'lonCell' in 'mesh' pool
------------	--------------------------

Table B.29: lonCell: Longitude location of cell centers in radians.

B.1.30 lonEdge

Type:	real
Units:	<i>radians</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'lonEdge' in 'mesh' pool

Table B.30: lonEdge: Longitude location of edge midpoints in radians.

B.1.31 lonVertex

Type:	real
Units:	<i>radians</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'lonVertex' in 'mesh' pool

Table B.31: lonVertex: Longitude location of vertices in radians.

B.1.32 nEdgesOnCell

Type:	integer
Units:	<i>unitless</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'nEdgesOnCell' in 'mesh' pool

Table B.32: nEdgesOnCell: Number of edges that border each cell.

B.1.33 nEdgesOnEdge

Type:	integer
Units:	<i>unitless</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'nEdgesOnEdge' in 'mesh' pool

Table B.33: nEdgesOnEdge: Number of edges that surround each of the cells that straddle each edge. These edges are used to reconstruct the tangential velocities.

B.1.34 sfcMassBal

Type:	real
Units:	$kg\ m^2\ s^{-1}$
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	input restart
Pool path:	'sfcMassBal' in 'mesh' pool

Table B.34: sfcMassBal: Surface mass balance

B.1.35 verticesOnCell

Type:	integer
Units:	<i>unitless</i>
Dimension:	maxEdges nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'verticesOnCell' in 'mesh' pool

Table B.35: verticesOnCell: List of vertices that border each cell.

B.1.36 verticesOnEdge

Type:	integer
Units:	<i>unitless</i>
Dimension:	TWO nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'verticesOnEdge' in 'mesh' pool

Table B.36: verticesOnEdge: List of vertices that straddle each edge.

B.1.37 weightsOnEdge

Type:	real
Units:	<i>unitless</i>
Dimension:	maxEdges2 nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'weightsOnEdge' in 'mesh' pool

Table B.37: weightsOnEdge: Reconstruction weights associated with each of the edgesOnEdge.

B.1.38 xCell

Type:	real
Units:	<i>unitless</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'xCell' in 'mesh' pool

Table B.38: xCell: X Coordinate in cartesian space of cell centers.

B.1.39 xEdge

Type:	real
Units:	<i>unitless</i>
Dimension:	nEdges

Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'xEdge' in 'mesh' pool

Table B.39: xEdge: X Coordinate in cartesian space of edge midpoints.

B.1.40 xVertex

Type:	real
Units:	<i>unitless</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'xVertex' in 'mesh' pool

Table B.40: xVertex: X Coordinate in cartesian space of vertices.

B.1.41 yCell

Type:	real
Units:	<i>unitless</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'yCell' in 'mesh' pool

Table B.41: yCell: Y Coordinate in cartesian space of cell centers.

B.1.42 yEdge

Type:	real
Units:	<i>unitless</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'yEdge' in 'mesh' pool

Table B.42: yEdge: Y Coordinate in cartesian space of edge midpoints.

B.1.43 yVertex

Type:	real
Units:	<i>unitless</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'yVertex' in 'mesh' pool

Table B.43: yVertex: Y Coordinate in cartesian space of vertices.

B.1.44 zCell

Type:	real
Units:	<i>unitless</i>
Dimension:	nCells
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'zCell' in 'mesh' pool

Table B.44: zCell: Z Coordinate in cartesian space of cell centers.

B.1.45 zEdge

Type:	real
Units:	<i>unitless</i>
Dimension:	nEdges
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'zEdge' in 'mesh' pool

Table B.45: zEdge: Z Coordinate in cartesian space of edge midpoints.

B.1.46 zVertex

Type:	real
Units:	<i>unitless</i>
Dimension:	nVertices
Persistence:	persistent
Number of time levels:	1
In streams:	basicmesh
Pool path:	'zVertex' in 'mesh' pool

Table B.46: zVertex: Z Coordinate in cartesian space of vertices.

B.2 state

B.2.1 cellMask

Type:	integer
Units:	<i>none</i>
Dimension:	nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	restart output
Pool path:	'cellMask' in 'state' pool

Table B.47: cellMask: bitmask indicating various properties about the ice sheet on cells.
cellMask only needs to be a restart field if config_allow_additional_advance = false (to keep the mask of initial ice extent)

B.2.2 edgeMask

Type:	integer
Units:	<i>none</i>
Dimension:	nEdges Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'edgeMask' in 'state' pool

Table B.48: edgeMask: bitmask indicating various properties about the ice sheet on edges.

B.2.3 layerThickness

Type:	real
Units:	m
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'layerThickness' in 'state' pool

Table B.49: layerThickness: layer thickness

B.2.4 layerThicknessEdge

Type:	real
Units:	m
Dimension:	nVertLevels nEdges Time
Persistence:	persistent
Number of time levels:	2
Pool path:	'layerThicknessEdge' in 'state' pool

Table B.50: layerThicknessEdge: layer thickness on cell edges

B.2.5 lowerSurface

Type:	real
Units:	m above datum
Dimension:	nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'lowerSurface' in 'state' pool

Table B.51: lowerSurface: elevation at bottom of ice

B.2.6 normalVelocity

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nEdges Time

Persistence:	persistent
Number of time levels:	2
In streams:	input restart output
Pool path:	'normalVelocity' in 'state' pool

Table B.52: normalVelocity: horizontal velocity, normal component to an edge

B.2.7 temperature

Type:	real
Units:	<i>degrees Celsius</i>
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
Index in temperature Array:	'temperature' in 'state' pool
Pool path:	'tracers' in 'state' pool
Array Group:	dynamics

Table B.53: temperature: ice temperature

B.2.8 thickness

Type:	real
Units:	<i>m</i>
Dimension:	nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	input restart output
Pool path:	'thickness' in 'state' pool

Table B.54: thickness: ice thickness

B.2.9 uReconstructMeridional

Type:	real
Units:	<i>m s⁻¹</i>
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output

Pool path:	'uReconstructMeridional' in 'state' pool
------------	--

Table B.55: uReconstructMeridional: meridional velocity reconstructed on cell centers

B.2.10 uReconstructX

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'uReconstructX' in 'state' pool

Table B.56: uReconstructX: x-component of velocity reconstructed on cell centers

B.2.11 uReconstructY

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'uReconstructY' in 'state' pool

Table B.57: uReconstructY: y-component of velocity reconstructed on cell centers

B.2.12 uReconstructZ

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'uReconstructZ' in 'state' pool

Table B.58: uReconstructZ: z-component of velocity reconstructed on cell centers

B.2.13 uReconstructZonal

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'uReconstructZonal' in 'state' pool

Table B.59: uReconstructZonal: zonal velocity reconstructed on cell centers

B.2.14 upperSurface

Type:	real
Units:	$m\ above\ datum$
Dimension:	nCells Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'upperSurface' in 'state' pool

Table B.60: upperSurface: elevation at top of ice

B.2.15 upperSurfaceVertex

Type:	real
Units:	$m\ above\ datum$
Dimension:	nVertices Time
Persistence:	persistent
Number of time levels:	2
Pool path:	'upperSurfaceVertex' in 'state' pool

Table B.61: upperSurfaceVertex: elevation at top of ice on vertices (currently only needed by shallow ice solver)

B.2.16 vertexMask

Type:	integer
Units:	<i>none</i>
Dimension:	nVertices Time
Persistence:	persistent
Number of time levels:	2
In streams:	output
Pool path:	'vertexMask' in 'state' pool

Table B.62: vertexMask: bitmask indicating various properties about the ice sheet on vertices.

B.2.17 xtime

Type:	text
Units:	<i>unitless</i>
Dimension:	Time
Persistence:	persistent
Number of time levels:	1
In streams:	restart output
Pool path:	'xtime' in 'state' pool

Table B.63: xtime: model time, with format 'YYYY-MM-DD_HH:MM:SS'

B.3 tend

B.3.1 tend_layerThickness

Type:	real
Units:	$m\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	1
Pool path:	'tend_layerThickness' in 'tend' pool

Table B.64: tend_layerThickness: time tendency of layer thickness

B.3.2 tend_temperature

Type:	real
-------	------

Units:	$K\ s^{-1}$
Dimension:	nVertLevels nCells Time
Persistence:	persistent
Number of time levels:	1
Index in tend_temperature Array:	'tend_temperature' in 'tend' pool
Pool path:	'tendTracers' in 'tend' pool
Array Group:	dynamics

Table B.65: tend_temperature: time tendency of ice temperature